



22883

PATENT TRADEMARK OFFICE

211.1001.01

This application is filed in the name of the following inventors:

<i>Inventor Name</i>	<i>Citizenship</i>	<i>Residence City and State</i>
Pankaj GUPTA	India	Mountain View, California
Srinivasan VENKATACHARY	India	Sunnyvale, California

The assignee is *Sahasra Networks, Inc.*, a Delaware corporation having a place of business at 3045 Park Blvd., Palo Alto, California.

TITLE OF THE INVENTION

High-Speed Message Forwarding Lookups for Arbitrary Length Strings Using Pipelined Memories

BACKGROUND OF THE INVENTION

1. *Field of the Invention*

This invention relates to message forwarding, such as for example in a computer network, wireless network or intra or inter computer communication system.

2. *Related Art*

In interconnection networks, individual devices communicate with each other using messages, which are forwarded using a collection of routers (or switches) from a source device to a destination device. Messages include header information, which is used by the routers to determine how to treat those messages. For example, each router must decide which next-hop router to forward a message to so that the message ultimately reaches its proper destination.

There are at least three concerns each router must address when forwarding messages. First, routers perform a lookup for routing information in response to packet header information. Second, routers schedule the packet for sending in response to a result of the lookup operation. Third, routers forward the packet to an output queue for sending, and actually send the packet to a next hop location.

When performing the lookup operation, there are further concerns each router must address. First, routers need a rapid technique for determining how to treat the message in response to the destination address, including obtaining the “longest prefix match” for the header information. Second, routers need an efficient technique for recording all known destination prefixes addresses in memory and for updating that memory in response to changes in network topology. (While this operation does not take place for each lookup, the router still performs the operation from time to time.) Given the so-

1 far rapid growth of the Internet, and the consequent rapid growth in both the number of
2 possible destinations and the number of routing table entries, every step in this process
3 requires accuracy and would benefit from speed.

4
5 The “longest prefix match” problem may be shown by an example. The
6 destination address of a message directed to MIT maybe routed a first way; such mes-
7 sages may include a prefix that is broad and general. The destination address of a mes-
8 sage directed to MIT’s School of Architecture may be routed a second way, different
9 from the first way and having priority over the first way. These distinctions and priorities
10 are made by looking up the destination address in a table that associates a destination ad-
11 dress with a next hop router and give preference to table entries that match the longest
12 part of the destination address (this is referred to as “longest match”). Given the rapid ex-
13 pansion of the Internet, tables are relatively large and growing relatively rapidly. Hence,
14 identifying the “longest match” has become increasingly difficult.

15
16 Longest prefix matching can be performed using a TCAM (ternary content
17 addressable memory) or by using various different algorithms such as different trie
18 structures. Both TCAMs and the various different algorithms have numerous drawbacks.
19 TCAMS have relatively low memory density, perform at a limited clock speed, and so
20 require multiple chips in order to achieve either relatively high speeds or relatively large
21 table sizes.

Moreover, algorithms used in lookups have relatively large and elaborate data structures, thus requiring the use of off-chip memory (additional memory not integrated onto the same chip as the lookup circuits).

Known methods of matching destination addresses primarily use off-chip memory, in which the router (or switch) communicates with a memory device using a memory bus to retrieve information about routing messages. When off-chip memory is used, the memory can be either very dense but relatively slow, such as DRAM, or can be very fast but relatively less dense, such as off-chip SRAM. Using DRAM, the speed with which the router can reference memory (and the amount of bandwidth for memory access) limits the speed at which the router can process header information. The relative lesser density of off-chip SRAM means that the router uses multiple SRAM devices for the same amount of storage, and therefore devotes a relatively large number of its input/output pins to accessing memory (also limiting the ability to quickly access large amounts of memory). Either of these drawbacks severely limits the speed with which the router can operate.

Alternatively, routers may use “on-chip” memory, in which the router includes memory with routing information integrated onto the same monolithic semiconductor circuit with intelligent logic. However, on-chip SRAM memory sizes are relatively limited, so that relatively compact or small data structures are preferred. On-chip

1 DRAM memory sizes are less limited, but on-chip DRAM is comparatively slow, and
2 therefore less suitable for rapid lookup.

3
4 TCAMs have the advantage of not requiring large and elaborate data struc-
5 tures, because they are specially designed for the purpose of lookup, but they have rela-
6 tively high power usage, low speed and low density. Avoiding specialized memories
7 leads to using other data structures, which in the known art are relatively large and there-
8 fore use multiple memory devices—either multiple SRAM chips (to achieve larger stor-
9 age), or multiple DRAM chips (to achieve higher speed).

10
11 Accordingly, it would be desirable to provide a technique for lookup of
12 message header information that is not subject to drawbacks of the known art. This can
13 be achieved using aspects of the invention in which a relatively compact and smaller data
14 structure is used in combination with pipelined on-chip memory.

15 16 SUMMARY OF THE INVENTION

17
18 The invention provides a method and system for lookup of message header
19 information that has the advantages of (1) the relative low scaling factor for power usage
20 associated with the use of a random access memory such as SRAM, which has relatively
21 constant power requirement relative to the amount of stored data, over that of associative
22 memories such as TCAM, which has power requirement proportional to the amount of

1 stored data, (2) the relative flexibility, associated with the use of a random access mem-
2 ory such as SRAM, over that of more specialized memories such as TCAM, (2) the rela-
3 tive speed associated with on-chip memory over that of off-chip memory or associative
4 memory. The invention includes methods and systems for lookup using pipelined on-chip
5 memory to access any relatively compact and smaller data structures stored on-chip; one
6 such relatively compact data structure is described in the incorporated disclosure, as re-
7 ferred to herein.

8
9 In a first aspect of the invention, lookups are performed using a sequence of
10 pipelined on-chip memories, each having only a portion of the header information in-
11 tended for lookup. Each one of the sequence of on-chip memories simultaneously per-
12 forms a lookup on a portion of the header information, thus allowing embodiments of the
13 invention to operate on multiple messages worth of header information substantially si-
14 multaneously. Even though each one lookup might take the full sequence of on-chip
15 memories to complete, the router is still able to receive and to process one message for
16 each memory cycle. Lookup speed is limited only by on-chip memory speed, rather than
17 by additional limitations of the known art. Given that on-chip memory is the fastest of all
18 available types of memory, lookup speed (and therefore router capability) scales well
19 with raw technology. Moreover, the use of a pipelined on-chip memory is particularly
20 cost effective.

1 The invention has general applicability to compact memory storage and re-
2 trieval, to update of information recorded in relatively compact form, and to applications
3 having utility for data lookup. None of these applications are limited specifically to
4 lookup for message header information, nor are they necessarily related to the specific
5 applications disclosed herein. For example, embodiments of the invention can include
6 one or more of, or some combination of, the following applications:

- 8 • Routers and switches at all levels of the ISO/OSI networking model, including
9 without limitation: bridges, network gateways, network routers and switches.
- 11 • Emulation of an associative memory with contiguous prefix masking, including
12 applications other than routing or switching.
- 14 • Both fixed-length and variable-length lookups in various types of lookup tables,
15 such as for example for different protocols such as IP, Ethernet, MPLS, and differ-
16 ent message types, such as for example unicast, multicast, and the like.

18 BRIEF DESCRIPTION OF THE DRAWINGS

19
20 Figure 1 shows a block diagram of a system for pipelined memory organi-
21 zation for flexible prefix matching.

Figure 2 shows a process flow diagram of a method for storage and retrieval of information using a system for pipelined memory organization for flexible prefix matching.

DESCRIPTION OF THE PREFERRED EMBODIMENT

The invention is described herein with regard to preferred steps and data structures. Those skilled in the art will recognize, after perusal of this application, that the described steps and data structures are not limited to any particular circuits. Rather, those of ordinary skill in the art would be able to implement the described steps and data structures, and equivalents thereof, without undue experimentation or further invention. All such implementations are within the scope and spirit of the invention.

Related Applications

Inventions described herein can be used in conjunction with inventions described in the following documents:

- U.S. Patent Application Serial No. _____, Express Mail Mailing No. EL734816260US, filed the same day in the name of the same inventors, attorney docket number 211.1002.01, titled "Compact Data Structures for Pipelined Message Forwarding Lookups".

1
2 This document is hereby incorporated by reference as if fully set forth
3 herein. This document is referred to as the incorporated disclosure.
4

5 A preferred embodiment includes data structures and memories as further
6 described in the incorporated disclosure.
7

8 *Lexicography*

9
10 The following terms refer or relate to aspects of the invention as described
11 below. The descriptions of general meanings of these terms are not intended to be limit-
12 ing, only illustrative.

- 13
14 • **router** — in general, any device for performing lookup of message header infor-
15 mation against an information database, and for treating messages having that
16 message header information differentially in response to a result of the lookup. For
17 example, a router can act in response to destination IP addresses (for forwarding),
18 in response to destination IP addresses and source IP addresses (for multicast for-
19 warding and for access control), or in response to other packet header fields (for
20 enforcing administrative restrictions or other message routing rules).

1
2 As noted above, these descriptions of general meanings of these terms are
3 not intended to be limiting, only illustrative. Other and further applications of the inven-
4 tion, including extensions of these terms and concepts, would be clear to those of ordi-
5 nary skill in the art after perusing this application. These other and further applications
6 are part of the scope and spirit of the invention, and would be clear to those of ordinary
7 skill in the art, without further invention or undue experimentation.

8
9 *System Elements*

10
11 Figure 1 shows a block diagram of a system for pipelined memory organi-
12 zation for flexible prefix matching.

13
14 A system 100 includes a lookup circuit 101, and can include an (optional)
15 off-chip memory 105, coupled to the lookup circuit 101 using a bus 106 (such as for ex-
16 ample a memory bus).

17
18 As noted above, the off-chip memory 105 and the bus 106 are optional. A
19 preferred embodiment does not require them, but can include them if desired to provide
20 additional memory accessible by the lookup circuit 101.

1 In a preferred embodiment, the lookup circuit 101 includes a single mono-
2 lithic circuit integrated onto at least one side of a silicon wafer (“chip”). Thus, the lookup
3 circuit 101 is preferably integrated onto a single chip, and the sequence of SRAM memo-
4 ries (as further described below) is thus accessible by techniques for accessing “on-chip”
5 memory.

6
7 The lookup circuit 101 includes an input port 102, an (optional) pre-lookup
8 circuit 103, a sequence of SRAM memories (as further described below), and an output
9 port 104.

10
11 There is no particular requirement that the sequence of memories must be
12 SRAM; they may be DRAM or may other type of on-chip memory or on-chip storage.
13 References to SRAM memory herein should be broadly read to encompass any method of
14 on-chip storage, now known or discovered in the future.

15
16 In a preferred embodiment, the input port 102 can be coupled to the bus
17 106, so as to be included in a larger system in which packet header information is sent to
18 the lookup circuit 101. However, in alternative embodiments, the input port 102 can be
19 directly coupled to an output of a circuit that delivers packet header information, or con-
20 figured to receive packet header information (or other information) in some other way.

1 The (optional) pre-lookup circuit 103 includes a set of pre-lookup logic 115
 2 and a pre-lookup memory circuit 120. Information coupled to the input port 102 is cou-
 3 pled to the pre-lookup logic 115 and the pre-lookup memory circuit 120. The pre-lookup
 4 memory circuit 120, under control of the pre-lookup logic 115, performs a direct memory
 5 lookup for prefix values that are relatively short (that is, smaller than a value d_0 , a se-
 6 lected number of bits). In embodiments of the invention, d_0 can be any value, so long as
 7 that value is sufficiently small to allow the lookup circuit 101 to operate effectively (for
 8 example, the pre-lookup memory circuit 120 must be fit onto the chip, and there may be
 9 other constraints). In a preferred embodiment, $d_0 = 14$, so prefix values of 14 or fewer
 10 bits are all recorded for lookup by the pre-lookup memory circuit 120, which can include
 11 a 2^{14} -entry table.

12
 13 The packet header information, having a total bit width of w , is coupled to
 14 the sequence of SRAM memories. Each stage n in the sequence includes an on-chip
 15 SRAM memory, an address lookup circuit, and a pipeline register, and spans d_n bits of
 16 the total bit width. Thus, a first stage in the sequence includes a first on-chip SRAM
 17 memory 125, a first address lookup circuit 130, and a first pipeline register 135, and
 18 spans d_1 bits of the total bit width; a second stage in the sequence includes a second on-
 19 chip SRAM memory 140, a second address lookup circuit 145, and a second pipeline
 20 register 150, and spans d_2 bits of the total bit width; and so on until a final stage in the
 21 sequence includes a final on-chip SRAM memory and a final address lookup circuit 160

1 (except as described below, there is no particular requirement for a final pipeline regis-
2 ter).

3
4 In a preferred embodiment, each value d_n is selected for optimal process-
5 ing of a lookup search key. In alternative embodiments, each value d_n may be distinct or
6 identical, may be selected to optimize a particular type of lookup, or may be arbitrarily
7 chosen.

8
9 In a preferred embodiment, while each SRAM memory has only a selected
10 number of entries, those entries can be quite wide in terms of the number of bits of in-
11 formation they provide, within only the limits of integrating SRAM memories onto a
12 chip. Thus for example, the result provided by an individual entry within one of the
13 SRAM memories can be a data structure having 256 bits, 1K bits, or more, depending on
14 the nature of the data stored therein and the physical size of the SRAM memories as inte-
15 grated onto the chip.

16
17 A result of the lookup operation is coupled to the corresponding stage n ad-
18 dress lookup circuit for interpretation. As described below, the lookup operation can pro-
19 vide a final lookup result (or a portion thereof) when there is a prefix entry for a prefix
20 length within the stride of d_n bits within the range of the corresponding stage n SRAM
21 memory. Also as described below, the lookup operation can provide a further lookup in-

1 dex into the corresponding next stage $n+1$ SRAM memory. The address lookup circuit
2 determines from the stage n lookup result whether it includes a final lookup result or a
3 further lookup index.

4
5 There is no particular requirement that the further lookup index must point
6 into the corresponding next stage $n+1$ SRAM memory. In some data structures, the fur-
7 ther lookup index may point into a later stage on-chip memory (a “jumping pointer” for-
8 ward into the sequence of memories), or may point into an earlier stage on-chip memory
9 (a “backward pointer”, particularly for loopback or multiple-loopback operation).

10
11 An output of the address lookup circuit is coupled to the corresponding
12 stage n pipeline register; the pipeline registers are clocked (preferably using a common
13 clock) so that the lookup operation can proceed at each stage n substantially in parallel.
14 Thus, as described herein, lookup #1 is being partially performed at stage n , lookup #2 is
15 being partially performed at stage $n-1$, ..., lookup # $n-1$ is being performed at stage 2,
16 and lookup # n is being performed at stage 1, all substantially simultaneously.

17
18 When a lookup is completed, the final lookup result is coupled to the output
19 port 104.

1 *Loopback Functionality*

2
3 In a first set of alternative embodiments, the lookup search key may be
4 longer than the width of the sequence of memories, and any remaining portion of the
5 lookup search key not yet referenced may be looped back to an earlier part (such as for
6 example the beginning) of the sequence. Thus, for example, where the packet header in-
7 formation has a total bit width of $2w$, the second half can be looped back to the beginning
8 of the sequence of memories, for further processing.

9
10 Loopback can be performed multiple times. Where the packet header in-
11 formation has a total bit width of $3w$, the second third can be looped back to the begin-
12 ning of the sequence of memories for further processing, and the final third can be looped
13 back a second time to the beginning of the sequence of memories for further processing.

14
15 In this first set of alternative embodiments, a result for the first w bits of the
16 lookup search key is maintained in the final pipeline register, and forwarded with the
17 portion of the lookup search key not yet processed (the second w bits in those cases when
18 the lookup search key is $2w$ bits wide, or the remaining $2w$ bits in those cases when the
19 lookup search key is $3w$ bits wide, and the like) to the stage 1 memory, or to a later stage
20 memory when the lookup search key is not an integer multiple of w bits.

1 Loopback need not be the very next lookup operation performed. In a pre-
2 ferred embodiment, any intermediate state is recorded with the loopback along with the
3 remaining part of the lookup search key, and the entire looped-back lookup operation is
4 queued with other lookup operations or update operations for being performed in a prior-
5 ity order. This prevents looped-back lookup operations from taking too much of the
6 memory lookup resources available.

7
8 This first set of alternative embodiments may be combined with the second
9 set of alternative embodiments, as described below. Thus, the lookup search key may be
10 any width; if more than w bits, the loopback functionality may be used to process the
11 portion of the lookup search key exceeding w bits, while if less than w bits, the multiple-
12 issue functionality (as described below) may be used to insert the lookup search key into
13 a stage n memory for n other than 1.

14 15 *Multiple-Issue Functionality*

16
17 In a second set of alternative embodiments, the lookup search key may be
18 inserted into the sequence of memories, and the lookup result may be provided from the
19 sequence of memories, at points other than the beginning and end of the sequence. This
20 has the advantages of (1) allowing for early reporting of lookup results where appropri-
21 ate; (2) allowing for parallel processing of multiple lookup search keys at once. Thus, for
22 example, if there are 24 stages in the sequence of memories, the first 8 stages may be

1 used for a first lookup, the second 8 stages may be used for a second lookup, and the third
2 8 stages may be used for a third lookup, all in parallel, thus providing a tripled lookup
3 rate.

4
5 In this second set of alternative embodiments, a set of multiple-issue logic
6 (not shown) inserts the lookup search key into the pipeline of memories at a stage n
7 memory (for n other than 1) and extracts the lookup result from the pipeline of memories
8 at a stage n memory (for n other than w). The multiple-issue logic determines, for each
9 lookup search key, where to insert the lookup search key and where to extract the lookup
10 result, so that parallel multiple issue of lookup results may be achieved for multiple
11 lookup search keys. In the example above, the multiple-issue logic would determine
12 which subsequence of the pipelined memories each lookup search key would be assigned
13 to.

14
15 As noted above, the first set of alternative embodiments may be combined
16 with the second set of alternative embodiments, so that portions of lookup search keys
17 less than w bits may be processed using multiple-issue functionality while portions of
18 lookup search keys when those lookup search keys exceed w bits may be processed using
19 loopback functionality.

1 *Method of Use*

2

3 Figure 2 shows a process flow diagram of a method for storage and re-
4 trieval of information using a system for pipelined memory organization for flexible pre-
5 fix matching.

6

7 A method 200 includes a set of flow points and process steps as described
8 herein.

9

10 Although by the nature of textual description, the flow points and process
11 steps are described sequentially, there is no particular requirement that the flow points or
12 process steps must be sequential. Rather, in preferred embodiments of the invention, the
13 described flow points and process steps are performed in a parallel or pipelined manner.

14

15 At a flow point 210, the lookup circuit 101 is ready to receive a lookup
16 search key at the input port 102.

17

18 At a step 211, the lookup circuit 101 couples the lookup search key to the
19 pre-lookup circuit 103.

20

21 In a preferred embodiment, the lookup circuit 101 also couples an instruc-
22 tion, such as a set of flag bits (not shown) indicating whether the lookup operation is in-

1 tended to store information in, or retrieve information from, the pipelined memory or-
2 ganization. Although this functionality is described herein as “storing” information, in a
3 preferred embodiment the flag bits can indicate any kind of modify or update operation,
4 such as to clear (that is, delete) an entry, mark an entry invalid, modify an entry to in-
5 clude new packet forwarding information, or otherwise alter the data structure maintained
6 by the sequence of pipelined on-chip memories.

7
8 At a step 212, the pre-lookup circuit 103 performs any pre-lookup opera-
9 tions. As noted above, the pre-lookup memory circuit 120, under control of the pre-
10 lookup logic 115, performs a direct memory lookup in response to the lookup search key.
11 As part of this step, the pre-lookup circuit 103 performs the following sub-steps:

- 12
- 13 • At a sub-step 212(a), the pre-lookup circuit 103 identifies the first d0 bits of the
14 lookup search key.
 - 15
 - 16 • At a sub-step 212(b), the pre-lookup circuit 103 performs a direct memory lookup
17 to search for matching prefix values less than d0 bits in length.
 - 18
 - 19 • At a sub-step 212(c), if the pre-lookup circuit 103 finds a matching prefix value, it
20 records that prefix value in a pre-lookup best-match register (not shown).
 - 21

1 As part of this sub-step, if the flag bits indicate that the lookup operation is to store
2 information, the pre-lookup circuit 103 stores any new information, such as a new
3 prefix value, in the identified location in the pre-lookup memory circuit 120.
4

5 At a step 213, the pre-lookup circuit 103 couples the lookup search key to
6 the sequence of SRAM memories.
7

8 At a step 214, each stage of the sequence of SRAM memories, in parallel,
9 performs a lookup on its portion of the lookup search key. As part of this step, the se-
10 quence of SRAM memories performs the following sub-steps:
11

- 12 • At a sub-step 214(a), each stage n in the sequence of SRAM memories identifies a
13 span of d_n bits of the lookup search key. As noted above, the number of bits d_n
14 might be different for each stage n in the sequence.
15
- 16 • At a sub-step 214(b), the n th on-chip SRAM memory performs a memory lookup
17 for a portion of a prefix value, in response to at least a portion of the identified
18 span of d_n bits.
19
- 20 • At a sub-step 214(c), the n th on-chip SRAM memory couples its lookup result to
21 the corresponding n th address lookup circuit.

1
2 As part of this sub-step, if the flag bits indicate that the lookup operation is to store
3 information, and the information to be stored is destined for a portion of the data
4 structure in the n th on-chip SRAM memory, that memory stores any new infor-
5 mation, such as a new prefix value, in the identified location therein.

6
7 In a preferred embodiment, as described in the incorporated disclosure,
8 more than one stage of the on-chip SRAM memory can cooperate in response to a
9 jointly-identified span of bits, to provide a single lookup result at the last one of those
10 more than one stages.

11
12 Accordingly, as described in the incorporated disclosure, if the flag bits in-
13 dicate that the lookup operation is to store information, more than one stage of the on-
14 chip SRAM memory can cooperate in response to a jointly-identified span of bits, to
15 store any new information, such as a new prefix value, in the identified locations in those
16 stages of the on-chip SRAM memory.

17
18 At a step 215, the corresponding n th address lookup circuit processes the
19 lookup result from the corresponding n th on-chip SRAM memory. As part of this step,
20 the corresponding n th address lookup circuit determines if the lookup result is either a fi-
21 nal lookup result (or a portion thereof), or an index into a corresponding stage $n+1$ in the

1 sequence of SRAM memories. If the lookup result is a final lookup result (or a portion
2 thereof), the corresponding nth address lookup circuit records the final lookup result (or
3 portion thereof) in its corresponding nth best-match register (part of the corresponding
4 nth pipeline register).

5
6 At step 216, the corresponding nth address lookup circuit couples its output
7 to the corresponding stage n pipeline register. As noted above, the pipeline registers are
8 clocked (preferably using a common clock) so that the lookup operation can proceed at
9 each stage n substantially in parallel.

10
11 At a flow point 220, the method 200 is complete, and the final lookup result
12 (if any) is coupled from the final stage best-match register 198 to the output port 104.

13
14 As described above, as part of the method, if the flag bits indicate that the
15 operation was to store (or otherwise modify or update) information in the data structure
16 maintained by the pipelined memories, a result of the method 200 is to store information
17 in the data structure synchronously with regard to lookup requests. Thus, each lookup re-
18 quest is performed with regard to a consistent data structure, because storage operations
19 occur explicitly either entirely before or entirely after the lookup request is performed,
20 even though those lookup operations and storage operations are performed substantially
21 concurrently by the pipelined memories.

Changes to the data structure that require more than one individual pass through the sequence of pipelined memory stages are performed (first) by adding any necessary new entries to each memory, and then (second) by adding or altering a pointer that links the new entries to the main data structure accessible by lookup operations. Addition of the final pointer is an atomic act as viewed by lookup operations. Thus, the main data structure, as viewed by lookup operations, is always consistent and is always updated atomically.

As described with regard to additional functionality above, it may occur that the final lookup result is coupled, along with the lookup search key, back to an earlier stage in the sequence of memories for further processing.

Generality of the Invention

The invention has general applicability to compact memory storage and retrieval, to update of information recorded in relatively compact form, and to applications having utility for data lookup. None of these applications are limited specifically to lookup for message header information, nor are they necessarily related to the specific applications disclosed herein. For example, embodiments of the invention can include one or more of, or some combination of, the following applications:

- 1 • Routers and switches at all levels of the ISO/OSI networking model, including
2 without limitation: bridges, network gateways, network routers and switches.
3
- 4 • Emulation of an associative memory with contiguous prefix masking, including
5 applications other than routing or switching.
6
- 7 • Both fixed-length and variable-length lookups in various types of lookup tables,
8 such as for example for different protocols such as IP, Ethernet, MPLS, and differ-
9 ent message types, such as for example unicast, multicast, and the like.
10

11 Moreover, techniques used by a preferred embodiment of the invention for
12 lookup of message header information can be used in contexts other than the specific ap-
13 plications disclosed herein. For example, techniques used by embodiments of the inven-
14 tion for storage and retrieval of information in relatively compact form, and to relatively
15 rapid pipelined data lookup, are all generally applicable to fields other than the specific
16 applications disclosed herein.
17

18 Other and further applications of the invention in its most general form
19 would be clear to those skilled in the art after perusal of this application. The invention
20 would be usable for such other and further applications without undue experimentation or
21 further invention.
22

1 Although preferred embodiments are disclosed herein, many variations are
2 possible which remain within the concept, scope and spirit of the invention; these varia-
3 tions would be clear to those skilled in the art after perusal of this application.

Express Mailing EL734816392US